IEEE Xplore® RELEASE 1.4

Welcome
United States Patent and Trademark Office

Your search matched **[0]** of **[945745]** documents.

You may refine your search by editing the current search expression or entering a new one the text box. Then click search Again.

cache ownership Search Again

OR

Use your browser's back button to return to your original search page.

Results:

No documents matched your query.

Terms used **multiprocessor** and **caches** and **global ownership** and **list**

Found **10,865** of **111,041**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Cache inclusion and processor sampling in multiprocessor simulations

Jacqueline Chame, Michel Dubois

June 1993 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems**, Volume 21 Issue 1

Full text available: pdf(1.17 MB) Additional Information: full citation, references, citings, index terms, review

2 Cache coherence in large-scale shared-memory multiprocessors: issues and comparisons

David J. Lilja

September 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 3

Full text available: pdf(3.12 MB) Additional Information: full citation, references, citings, index terms

3 SM-prof: a tool to visualise and find cache coherence performance bottlenecks in multiprocessor programs

Mats Brorsson

May 1995 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1995 ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems**, Volume 23 Issue 1

Full text available: pdf(1.58 MB) Additional Information: full citation, abstract, references, index terms

Cache misses due to coherence actions are often the major source for performance degradation in cache coherent multiprocessors. It is often difficult for the programmer to take cache coherence into account when writing the program since the resulting access pattern is not apparent until the program is executed.SM-prof is a performance analysis tool that addresses this problem by visualising the shared data access pattern in a diagram with links to the source code lines causing performance degrad ...

4 Flexible use of memory for replication/migration in cache-coherent DSM multiprocessors

Vijayaraghavan Soundararajan, Mark Heinrich, Ben Verghese, Kourosh Gharachorloo, Anoop Gupta, John Hennessy

April 1998 **ACM SIGARCH C mputer Architecture News , Pr ceedings f the 25th annual internati nal symposium n Computer architecture**, Volume 26 Issue 3

Full text available: pdf(1.76 MB) Additional Information: full citation, abstract, references, citings, index

terms

Given the limitations of bus-based multiprocessors, CC-NUMA is the scalable architecture of choice for shared-memory machines. The most important characteristic of the CC-NUMA architecture is that the latency to access data on a remote node is considerably larger than the latency to access local memory. On such machines, good data locality can reduce memory stall time and is therefore a critical factor in application performance.In this paper we study the various options available to system desi ...

5 The performance of cache-coherent ring-based multiprocessors

Luis André Barroso, Michel Dubois

May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture**, Volume 21 Issue 2

Full text available: pdf(1.03 MB) Additional Information: full citation, abstract, references, citings, index terms

Advances in circuit and integration technology are continuously boosting the speed of microprocessors. One of the main challenges presented by such developments is the effective use of powerful microprocessors in shared memory multiprocessor configurations. We believe that the interconnection problem is not solved even for small scale shared memory multiprocessors, since the speed of shared buses is unlikely to keep up with the bandwidth requirements of new microprocessors. In this paper we ...

6 Implications of hierarchical N-body methods for multiprocessor architectures

Jaswinder Pal Singh, John L. Hennessy, Anoop Gupta

May 1995 **ACM Transactions on Computer Systems (TOCS)**, Volume 13 Issue 2

Full text available: pdf(4.66 MB) Additional Information: full citation, abstract, references, citings, index terms, review

To design effective large-scale multiprocessors, designers need to understand the characteristics of the applications that will use the machines. Application characteristics of particular interest include the amount of communication relative to computation, the structure of the communication, and the local cache and memory requirements, as well as how these characteristics scale with larger problems and machines. One important class of applications is based on hierarchical N-body methods, w ...

Keywords: N-body methods, communication abstractions, locality, message passing, parallel applications, parallel computer architecture, scaling, shared address space, shared memory

7 Measuring memory hierarchy performance of cache-coherent multiprocessors using micro benchmarks

Cristina Hristea, Daniel Lenoski, John Keen

November 1997 **Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available: pdf(97.47 KB) Additional Information: full citation, abstract, references, citings

Even with today's large caches, the increasing performance gap between processors and memory systems imposes a memory bottleneck for many important scientific and commercial applications. This bottleneck is intensified in shared-memory multiprocessors by contention and the effects of cache coherency. Under heavy memory contention, the memory latency may increase 2 or 3 times. Nonethless, as more sophisticated techniques are used to hide latency and increase bandwidth, measuring memory performanc ...

8 The directory-based cache coherence protocol for the DASH multiprocessor

Daniel Lenoski, James Laudon, Kourosh Gharachorloo, Anoop Gupta, John Hennessy

May 1990 **ACM SIGARCH Computer Architecture News , Pr ceedings f the 17th annual international symposium on Computer Architecture**, Volume 18 Issue 3

Full text available: pdf(1.74 MB) Additional Information: full citation, abstract, references, citings, index terms

DASH is a scalable shared-memory multiprocessor currently being developed at Stanford's Computer Systems Laboratory. The architecture consists of powerful processing nodes, each with a portion of the shared-memory, connected to a scalable interconnection network. A

key feature of DASH is its distributed directory-based cache coherence protocol. Unlike traditional snoopy coherence protocols, the DASH protocol does not rely on broadcast; instead it uses point-to-point messages sent between th ...

9 Cache consistency in hierarchical-ring-based multiprocessors
K. Farkas, Z. Vranesic, M. Stumm
December 1992 **Proceedings of the 1992 ACM/IEEE conference on Supercomputing**

Full text available: pdf(906.50 KB) Additional Information: full citation, references, citings, index terms

10 Architecture: Leveraging cache coherence in active memory systems
Daehyun Kim, Mainak Chaudhuri, Mark Heinrich
June 2002 **Proceedings of the 16th international conference on Supercomputing**

Full text available: pdf(217.27 KB) Additional Information: full citation, abstract, references, index terms

Active memory systems help processors overcome the memory wall when applications exhibit poor cache behavior. They consist of either active memory elements that perform data parallel computations in the memory system itself, or an active memory controller that supports address re-mapping techniques that improve data locality. Both active memory approaches create coherence problems---even on uniprocessor systems---since there are either additional processors operating on the data directly, or the ...

Keywords: active memory, address re-mapping, cache coherence

11 Disk-directed I/O for MIMD multiprocessors
David Kotz
February 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 1

Full text available: pdf(559.18 KB) Additional Information: full citation, abstract, references, citings, index terms, review

Many scientific applications that run on today's multiprocessors, such as weather forecasting and seismic analysis, are bottlenecked by their file-I/O needs. Even if the multiprocessor is configured with sufficient I/O hardware, the file system software often fails to provide the available bandwidth to the application. Although libraries and enhanced file system interfaces can make a significant improvement, we believe that fundamental changes are needed in the file server software. We prop ...

Keywords: MIMD, collective I/O, disk-directed I/O, file caching, parallel I/O, parallel file system

12 Join and Semijoin Algorithms for a Multiprocessor Database Machine
Patrick Valduriez, Georges Gardarin
March 1984 **ACM Transactions on Database Systems (TODS)**, Volume 9 Issue 1

Full text available: pdf(1.95 MB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents and analyzes algorithms for computing joins and semijoins of relations in a multiprocessor database machine. First, a model of the multiprocessor architecture is described, incorporating parameters defining I/O, CPU, and message transmission times that permit calculation of the execution times of these algorithms. Then, three join algorithms are presented and compared. It is shown that, for a given configuration, each algorithm has an application domain defined by the ch ...

13 Architectural primitives for a scalable shared memory multiprocessor
Joonwon Lee, Umakishore Ramachandran
June 1991 **Proceedings of the third annual ACM symposium on Parallel algorithms and architectures**

Full text available: pdf(1.27 MB) Additional Information: full citation, references, citings, index terms

14 Cache coherence for large scale shared memory multiprocessors

M. Thapar, B. Delagi

May 1990 **Proceedings of the second annual ACM symposium n Parallel algorithms and architectures**

Full text available: pdf(645.67 KB) Additional Information: full citation, references, index terms

15 Data speculation support for a chip multiprocessor

Lance Hammond, Mark Willey, Kunle Olukotun

October 1998 **Proceedings of the eighth international conference on Architectural support for programming languages and operating systems**

Full text available: pdf(1.75 MB) Additional Information: full citation, abstract, references, citings, index terms

Thread-level speculation is a technique that enables parallel execution of sequential applications on a multiprocessor. This paper describes the complete implementation of the support for threadlevel speculation on the Hydra chip multiprocessor (CMP). The support consists of a number of software speculation control handlers and modifications to the shared secondary cache memory system of the CMP This support is evaluated using five representative integer applications. Our results show that the s ...

16 Verification techniques for cache coherence protocols

Fong Pong, Michel Dubois

March 1997 **ACM Computing Surveys (CSUR)**, Volume 29 Issue 1

Full text available: pdf(1.25 MB) Additional Information: full citation, abstract, references, citings, index terms

In this article we present a comprehensive survey of various approaches for the verification of cache coherence protocols based on state enumeration, (symbolic model checking, and symbolic state models. Since these techniques search the state space of the protocol exhaustively, the amount of memory required to manipulate that state information and the verification time grow very fast with the number of processors and the complexity of the protocol mechanism ...

Keywords: cache coherence, finite state machine, protocol verification, shared-memory multiprocessors, state representation and expansion

17 Characterizing the caching and synchronization performance of a multiprocessor operating system

Josep Torrellas, Anoop Gupta, John Hennessy

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems**, Volume 27 Issue 9

Full text available: pdf(1.52 MB) Additional Information: full citation, references, citings, index terms

18 Multiprocessor cache design considerations

R. L. Lee, P. C. Yew, D. H. Lawrie

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available: pdf(933.86 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper, cache design is explored for large high-performance multiprocessors with hundreds or thousands of processors and memory modules interconnected by a pipe-lined multi-stage network. The majority of the multiprocessor cache studies in the literature exclusively focus on the issue of cache coherence enforcement. However, there are other characteristics unique to such multiprocessors which create an environment for cache performance that is very different from that of many uniproc ...

19 A cache coherence scheme suitable for massively parallel processors
R. Baldwin
December 1993 **Proceedings of the 1993 ACM/IEEE conference on Supercomputing**
Full text available: pdf(931.96 KB) Additional Information: full citation, references, index terms

20 The implications of cache affinity on processor scheduling for multiprogrammed, shared memory multiprocessors
Raj Vaswani, John Zahorjan
September 1991 **ACM SIGOPS Operating Systems Review , Proceedings of the thirteenth ACM symposium on Operating systems principles**, Volume 25 Issue 5
Full text available: pdf(1.57 MB) Additional Information: full citation, references, citings, index terms

IEEE Xplore® RELEASE 1.4

Welcome
United States Patent and Trademark Office

Your search matched **7** of **942182** documents.

A maximum of **7** results are displayed, **15** to a page, sorted by **Relevance** in **descending** order.
You may refine your search by editing the current search expression or entering a new one the text box.
Then click **Search Again**.

Search Again

Results:
Journal or Magazine = **JNL** Conference = **CNF** Standard = **STD**

1 **A new cache directory scheme**
Yuguang Wu; Muntz, R.R.;
Parallel Architectures, Algorithms, and Networks, 1996. Proceedings. Second International Symposium on , 12-14 Jun 1996
Page(s): 466 -472

[Abstract] [PDF Full-Text (848 KB)] **IEEE CNF**

2 **Cooperative caching of dynamic content on a distributed Web server**
Holmedahl, V.; Smith, B.; Tao Yang;
High Performance Distributed Computing, 1998. Proceedings. The Seventh International Symposium on , 28-31 Jul 1998
Page(s): 243 -250

[Abstract] [PDF Full-Text (104 KB)] **IEEE CNF**

3 **Partial address directory for cache access**
Lishing Liu;
Very Large Scale Integration (VLSI) Systems, IEEE Transactions on , Volume: 2 Issue: 2 , Jun 1994
Page(s): 226 -240

[Abstract] [PDF Full-Text (1356 KB)] **IEEE JNL**

4 **Data structure distribution and multi-threading of Linux file system for multiprocessors**
Sheth, A.; Gopinath, K.;
High Performance Computing, 1998. HIPC '98. 5th International Conference On , 17-20 Dec 1998
Page(s): 97 -104

[Abstract] [PDF Full-Text (240 KB)] **IEEE CNF**

5 **A bandwidth friendly search engine**
Bradford, C.; Marshall, I.W.;
Multimedia Computing and Systems, 1999. IEEE International Conference on ,
Volume: 2 , Jul 1999
Page(s): 720 -724 vol.2

[Abstract] [PDF Full-Text (452 KB)] **IEEE CNF**

6 **Directory-based cache coherence in large-scale multiprocessors**
Chaiken, D.; Fields, C.; Kurihara, K.; Agarwal, A.;
Computer , Volume: 23 Issue: 6 , Jun 1990
Page(s): 49 -58

[Abstract] [PDF Full-Text (1164 KB)] **IEEE JNL**

7 **Summary cache: a scalable wide-area Web cache sharing protocol**
Li Fan; Pei Cao; Almeida, J.; Broder, A.Z.;
Networking, IEEE/ACM Transactions on , Volume: 8 Issue: 3 , Jun 2000
Page(s): 281 -293

[Abstract] [PDF Full-Text (220 KB)] **IEEE JNL**

Welcome to IEEE Xplore®

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

Print Format

Your search matched **3** of **942182** documents.

A maximum of **3** results are displayed, **15** to a page, sorted by **Relevance** in **descending** order.
You may refine your search by editing the current search expression or entering a new one the text box.
Then click **Search Again**.

directory "> Search Again

Results:
Journal or Magazine = **JNL** Conference = **CNF** Standard = **STD**

1 **A multi-cache coherence scheme for shuffle-exchange network based multiprocessors**
Omran, R.A.; De-Lei Lee;
Frontiers of Massively Parallel Computation, 1995. Proceedings. 'Frontiers '95'., Fifth Symposium on the , 6-9 Feb 1995
Page(s): 72 -79

[Abstract] [PDF Full-Text (740 KB)] **IEEE CNF**

2 **Reduce, reuse, recycle: an approach to building large Internet caches**
Gadde, S.; Rabinovich, M.; Chase, J.;
Operating Systems, 1997., The Sixth Workshop on Hot Topics in , 5-6 May 1997
Page(s): 93 -98

[Abstract] [PDF Full-Text (540 KB)] **IEEE CNF**

3 **Cooperative caching of dynamic content on a distributed Web server**
Holmedahl, V.; Smith, B.; Tao Yang;
High Performance Distributed Computing, 1998. Proceedings. The Seventh International Symposium on , 28-31 Jul 1998
Page(s): 243 -250

[Abstract] [PDF Full-Text (104 KB)] **IEEE CNF**

Terms used **cache directory** and **cache wnership**

Found **57** of **110,178**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 57

Result page: **1** 2 3 next

Relevance scale

1 Techniques for fast simulation of associative cache directories

Mayan Moudgill

May 1998 **ACM SIGARCH Computer Architecture News**, Volume 26 Issue 2

Full text available: pdf(527.92 KB) Additional Information: full citation, abstract, index terms

We describe a technique for simulating associative cache directories that considerably reduces simulation time with respect to a sequential search of the tag array. We also describe techniques for maintaining LRU information that use considerably less memory than time-stamps. The combination of these techniques makes possible the simulation of set-associative cache directories at a much higher speed than other techniques.

2 Cache designs with partial address matching

Lishing Liu

November 1994 **Proceedings of the 27th annual international symposium on Microarchitecture**

Full text available: pdf(938.27 KB) Additional Information: full citation, abstract, references, citings, index terms

One critical aspect in designing set-associative cache at high clock rate is deriving timely results from directory lookup. In this paper we investigate the possibility of accurately approximating the results of conventional directory search with faster matches of few partial address bits. Such fast and accurate approximations may be utilized to optimize cache access timing, particularly in a customized design environment. Through analytic and simulation studies we examine the trade-offs of ...

3 Research sessions: distributed systems: Proxy-based acceleration of dynamically generated content on the world wide web: an approach and implementation

Anindya Datta, Kaushik Dutta, Helen Thomas, Debra VanderMeer, Debra VanderMeer, Krithi Ramamritham

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available: pdf(1.37 MB) Additional Information: full citation, abstract, references, index terms

As Internet traffic continues to grow and web sites become increasingly complex, performance and scalability are major issues for web sites. Web sites are increasingly relying on dynamic content generation applications to provide web site visitors with dynamic, interactive, and personalized experiences. However, dynamic content generation comes at a cost --- each request requires computation as well as communication across multiple components.To address these issues, various dynamic content cach ...

Keywords: dynamic content, edge caching, proxy-based caching

4 The VMP multiprocessor: initial experience, refinements, and performance evaluation

D. R. Cheriton, A. Gupta, P. D. Boyle, H. A. Goosen

May 1988 **ACM SIGARCH Computer Architecture News , Pr ceedings of the 15th**

Annual International Symposium on Computer architecture, Volume 16 Issue 2

Full text available: pdf(1.73 MB) Additional Information: full citation, abstract, references, citings, index terms

VMP is an experimental multiprocessor being developed at Stanford University, suitable for high-performance workstations and server machines. Its primary novelty lies in the use of software management of the per-processor caches and the design decisions in the cache and bus that make this approach feasible. The design and some uniprocessor trace-driven simulations indicating its performance have been reported previously. In this paper, we present our initial experience with the V ...

5 Architectural primitives for a scalable shared memory multiprocessor

Joonwon Lee, Umakishore Ramachandran

June 1991 **Proceedings of the third annual ACM symposium on Parallel algorithms and architectures**

Full text available: pdf(1.27 MB) Additional Information: full citation, references, citings, index terms

6 Summary cache: a scalable wide-area web cache sharing protocol

Li Fan, Pei Cao, Jussara Almeida, Andrei Z. Broder

June 2000 **IEEE/ACM Transactions on Networking (TON)**, Volume 8 Issue 3

Full text available: pdf(220.29 KB) Additional Information: full citation, references, citings, index terms

Keywords: ICP, Web cache, Web proxy, bloom filter, cache sharing

7 An architectural perspective on a memory access controller

M. Freeman

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available: pdf(1.08 MB) Additional Information: full citation, abstract, references, citings, index terms

In this paper a CMOS memory access controller chip is described that provides the basis for achieving high-performance 68020-based (68030-based) systems. This controller matches the speed of the memory system to that of the microprocessor by providing a virtual cache mechanism where address translations are only required when there is a cache miss. This mechanism also facilitates the construction of shared-memory multiprocessor system where the controller manages ...

8 A low-overhead coherence solution for multiprocessors with private cache memories

Mark S. Papamarcos, Janak H. Patel

January 1984 **ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture**, Volume 12 Issue 3

Full text available: pdf(590.93 KB) Additional Information: full citation, abstract, references, citings, index terms

This paper presents a cache coherence solution for multiprocessors organized around a single time-shared bus. The solution aims at reducing bus traffic and hence bus wait time. This in turn increases the overall processor utilization. Unlike most traditional high-performance coherence solutions, this solution does not use any global tables. Furthermore, this coherence scheme is modular and easily extensible, requiring no modification of cache modules to add more processors to a system. The ...

9 Cache design of a sub-micron CMOS system/370

J. H. Chang, H. Chao, K. So

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available: pdf(466.38 KB) Additional Information: full citation, abstract, references, citings, index terms

An innovative cache accessing scheme based on high MRU (most recently used) hit ratio [1] is proposed for the design of a one-cycle cache in a CMOS implementation of System/370. It is shown that with this scheme the cache access time is reduced by 30 ~ 35% and the performance is within 4% of a true one-cycle cache. This cache scheme is proposed to be used in a VLSI System/370, which is organized to achieve high performance by taking advantage of the performance and integration level of an a ...

10 Increasing web server throughput with network interface data caching
Hyong-youb Kim, Vijay S. Pai, Scott Rixner
October 2002 **Tenth international conference on architectural support for programming languages and operating systems on Proceedings of the 10th international conference on architectural support for programming languages and operating systems (ASPLOS-X)**
Full text available: pdf(1.22 MB) Additional Information: full citation, abstract, references

This paper introduces network interface data caching, a new technique to reduce local interconnect traffic on networking servers by caching frequently-requested content on a programmable network interface. The operating system on the host CPU determines which data to store in the cache and for which packets it should use data from the cache. To facilitate data reuse across multiple packets and connections, the cache only stores application-level response content (such as HTTP data), with applica ...

11 An economical solution to the cache coherence problem
James Archibald, Jean Loup Baer
January 1984 **ACM SIGARCH Computer Architecture News , Proceedings of the 11th annual international symposium on Computer architecture**, Volume 12 Issue 3
Full text available: pdf(728.73 KB) Additional Information: full citation, abstract, references, citings, index terms

In this paper we review and qualitatively evaluate schemes to maintain cache coherence in tightly-coupled multiprocessor systems. This leads us to propose a more economical (hardware-wise), expandable and modular variation of the "global directory" approach. Protocols for this solution are described. Performance evaluation studies indicate the limits (number of processors, level of sharing) within which this approach is viable.

12 Cache coherence in large-scale shared-memory multiprocessors: issues and comparisons
David J. Lilja
September 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 3
Full text available: pdf(3.12 MB) Additional Information: full citation, references, citings, index terms

13 Using LDAP directory caches
Sophie Cluet, Olga Kapitskaia, Divesh Srivastava
May 1999 **Proceedings of the eighteenth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems**
Full text available: pdf(1.32 MB) Additional Information: full citation, references, citings, index terms

14 Implementing global memory management in a workstation cluster
M. J. Feeley, W. E. Morgan, E. P. Pighin, A. R. Karlin, H. M. Levy, C. A. Thekkath
December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5
Full text available: pdf(1.52 MB) Additional Information: full citation, references, citings, index terms

15 Effects of cache coherency in multiprocessors
Michel Dubois, Fayé A. Briggs
April 1982 **Proceedings of the 9th annual symposium on Computer Architecture**
Additional Information: full citation, abstract, references, citings, index

Full text available: pdf(799.09 KB) terms

In many commercial multiprocessor systems, each processor accesses the memory through a private cache. One problem that could limit the extensibility of the system and its performance is the enforcement of cache coherence. A mechanism must exist which prevents the existence of several different copies of the same data block in different private caches. In this paper, we present an indepth analysis of the effect of cache coherency in multiprocessors. A novel analytical model for the program ...

16 Cache coherence protocols: evaluation using a multiprocessor simulation model

James Archibald, Jean-Loup Baer

September 1986 **ACM Transactions on Computer Systems (TOCS)**, Volume 4 Issue 4

Full text available: pdf(1.79 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Using simulation, we examine the efficiency of several distributed, hardware-based solutions to the cache coherence problem in shared-bus multiprocessors. For each of the approaches, the associated protocol is outlined. The simulation model is described, and results from that model are presented. The magnitude of the potential performance difference between the various approaches indicates that the choice of coherence solution is very important in the design of an efficient shared-bus multi ...

17 RISCY patents

David A. Patterson

September 1988 **ACM SIGARCH Computer Architecture News**, Volume 16 Issue 4

Full text available: pdf(1.83 MB) Additional Information: full citation, index terms

18 A new cache replacement scheme based on backpropagation neural networks

Humayun Khalid

March 1997 **ACM SIGARCH Computer Architecture News**, Volume 25 Issue 1

Full text available: pdf(572.97 KB) Additional Information: full citation, abstract, index terms

In this paper, we present a new neural network-based algorithm, **KORA** (*K*halid Shad*O*w *R*eplacement *A*lgorithm), that uses backpropagation neural network (BPNN) for the purpose of guiding the line/block replacement decisions in cache. This work is a continuation of our previous research presented in [1]-[3]. The KORA algorithm attempts to approximate the replacement decisions made by the optimal scheme (OPT). The key to our algorithm is to identify and subsequently ...

Keywords: cache memory, neural networks, performance evaluation

19 Reducing the frequency of tag compares for low power I-cache design

Ramesh Panwar, David Rennels

April 1995 **Proceedings 1995 international symposium on Low power design**

Full text available: pdf(447.34 KB) Additional Information: full citation, references, citings, index terms

20 Recovery protocols for shared memory database systems

Lory D. Molesky, Krithi Ramamritham

May 1995 **ACM SIGMOD Record , Proceedings of the 1995 ACM SIGMOD international c nference on Management of data**, Volume 24 Issue 2

Full text available: pdf(1.65 MB) Additional Information: full citation, abstract, references, index terms

Significant performance advantages can be gained by implementing a database system on a cache-coherent shared memory multiprocessor. However, problems arise when failures occur. A single node (where a *node* refers to a processor/memory pair) crash may require a reboot of the entire shared memory system. Fortunately, shared memory multiprocessors that isolate individual node failures are currently being developed. Even with these, because of the side effects of the cache coherency protocol, ...

Terms used **global cache directory** and **cache ownership**

Found **1** of **110,178**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 1 of 1

Relevance scale

1 Implementing global memory management in a workstation cluster

M. J. Feeley, W. E. Morgan, E. P. Pighin, A. R. Karlin, H. M. Levy, C. A. Thekkath

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth ACM symposium on Operating systems principles**, Volume 29 Issue 5

Full text available: pdf(1.52 MB) Additional Information: full citation, references, citings, index terms

Results 1 - 1 of 1